

A stochastic epidemiological model and a deterministic limit for BitTorrent-like peer-to-peer file-sharing networks

GEORGE KESIDIS ^{*}
Pennsylvania State University

TAKIS KONSTANTOPOULOS [†]
Heriot-Watt University

PERLA SOUSI [‡]
University of Cambridge

Abstract

In this paper, we propose a stochastic model for a file-sharing peer-to-peer network which resembles the popular BitTorrent system: large files are split into chunks and a peer can download or swap from another peer only one chunk at a time. We prove that the fluid limits of a scaled Markov model of this system are of the coagulation form, special cases of which are well-known epidemiological (SIR) models. In addition, Lyapunov stability and settling-time results are explored. We derive conditions under which the BitTorrent incentives under consideration result in shorter mean file-acquisition times for peers compared to client-server (single chunk) systems. Finally, a diffusion approximation is given and some open questions are discussed.

Keywords and phrases. FLUID LIMIT, PEER-TO-PEER NETWORKS, DIFFUSION APPROXIMATION, STABILITY, EPIDEMIOLOGICAL MODEL

AMS 2000 subject classifications. Primary 60F17, 90B15; secondary 60J75

1 Introduction

Peer-to-peer (p2p) activity continues to represent a very significant fraction of overall Internet traffic, 44% by one recent account [4]. BitTorrent [1, 2, 8, 21, 18, 9, 19] is a widely deployed p2p file-sharing network which has recently played a significant role in the network neutrality debate. Under BitTorrent, peers join “swarms” (or “torrents”) where each swarm corresponds to a specific data object (file). The process of finding the peers in a given swarm to connect to is typically facilitated through a centralised “tracker”. Recently, a trackerless BitTorrent client has been introduced that uses distributed hashing for query resolution [16].

For file sharing, a peer is typically uploads upload pieces (“chunks”) of the file to other peers in the swarm while downloading his/her missing chunks from them. This chunk swapping constitutes a transaction-by-transaction incentive for peers to cooperate (i.e., trading

^{*}CSE and EE Depts, Pennsylvania State University, University Park, PA 16802, USA; gik2@psu.edu

[†]School of Mathematical & Computer Sciences and the Maxwell Institute for Mathematical Sciences, Heriot-Watt University, Edinburgh EH14 4AS, UK; takis@ma.hw.ac.uk

[‡]Statistical Laboratory, Centre for Mathematical Sciences, University of Cambridge, Wilberforce Road, Cambridge CB3 0WB, UK; P.Sousi@statslab.cam.ac.uk

rather than simply download) to disseminate data objects. Large files may be segmented into several hundred chunks, all of which the peers of the corresponding swarm must collect, and in the process disseminate their own chunks before they can reconstitute the desired file and possibly leave the file’s swarm.

In addition to the framework in which data objects are segmented into chunks to promote cooperation through swapping, there is a system whereby the rate at which chunks are uploaded is assessed for any given transaction, and peers that allocate inadequate bandwidth for uploading may be “choked” [14, 17]. Choking may also be applied to peers who, by employing multiple identities (sybils), abuse BitTorrent’s system of allowing newly arrived peers to a swarm to just download a few chunks (as they clearly cannot trade what they simply do not as yet possess). BitTorrent can also rehabilitate peers by (optimistically) unchoking them. In the following, we do not directly consider upload bandwidth and related choking issues.

In this paper, we motivate a deterministic epidemiological model of file dissemination for peer-to-peer file-sharing networks that employ BitTorrent-like incentives, a generalisation of that given in [10]*. Our model is different from those explored in [15, 21, 18] for BitTorrent, and we compute different quantities of interest. Our epidemiological framework, similar to that we used for the spread of multi-stage worms [12], could also be adapted for network coding systems. In [9], the authors propose a “fluid” model of a single torrent/swarm (as we do in the following) and fit it to (transient) data drawn from aggregate swarms. The connection to branching process models [21, 8] is simply that ours only tracks the number of active peers who possess or demand the file under consideration, i.e., a single swarm. Though our model is significantly simpler than that of prior work, it is derived directly from an intuitive transaction-by-transaction Markov process modelling file-dissemination of the p2p network and its numerical solutions clearly demonstrate the effectiveness of the aforementioned incentives. A basic assumption in the following is that peers do not distribute bogus files (or file chunks) [20].

This paper is organised as follows. The Markovian model is developed in detail in Section 2. A proof of its fluid limit is described in Section 3 (and in the Appendix) including Lyapunov stability and settling time results. The behaviour of the limiting ODE is studied in Section 4 for specific examples. In Section 5, we derive conditions under which the BitTorrent-like system has improved performance (smaller mean time to completely acquire the file by a peer) compared to a system of pure client-server (no chunk-swapping) interactions. A diffusion approximation is given in 6. The paper concludes with a discussion of open problems in Section 7.

2 The stochastic model

We fix a set F (a file) which is partitioned into n (on the order of hundreds) pieces called chunks. Consider a large networked “swarm” of N nodes called peers. Each peer possesses a certain (possibly empty) subset A of F . As time goes by, this peer interacts with other peers, the goal being to enlarge his set A until, eventually, the peer manages to collect all n chunks of F . The interaction between peers can either be a download or a swap; in both cases, chunks are being copied from peer to peer and are assumed never lost. A peer will stay in the network as long as he does not possess all chunks. After collecting everything,

*And this paper is a significant extension of [11].

sooner or later a peer departs or switches off. By splitting the desired file into many chunks we give incentives to the peers to remain active in the swarm for long time during which other peers will take advantage of their possessions.

2.1 Possible interactions

We here describe how two peers, labelled A, B , interact. The following types of interactions are possible:

1. **Download:** Peer A downloads a chunk i from B . This is possible only if A is a strict subset of B . If $i \in B$ then, after the downloading A becomes $A' = A \cup \{i\}$ and but B remains B because it since it gains nothing from A . Denote this interaction as:

$$\boxed{(A \leftarrow B) \rightsquigarrow (A', B)}$$

The symbol on the left is supposed to show the type of interaction and the labels before it, while the symbol on the right shows the labels after the interaction.

2. **Swap:** Peer A swaps with peer B . In other words, A gets a chunk j from B and B gets a chunk i from A . It is required that j is not an element of A and i not an element of B . We denote this interaction by

$$\boxed{(A \rightleftharpoons B) \rightsquigarrow (A', B')}$$

where $A' = A \cup \{j\}$, $B' = B \cup \{i\}$. We thus need $A \setminus B \neq \emptyset$ and $B \setminus A \neq \emptyset$.

2.2 Notation

The set of all combinations of n chunks, which partition F , is denoted by $\mathcal{P}(F)$, where $|\mathcal{P}(F)| = 2^n$ and the empty set is included. We write $A \subset B$ (respectively, $A \subsetneq B$) when A is a subset (respectively, strict subset) of B . We (unconventionally) write

$$A \sqsubset A' \text{ when } A \subset A' \text{ and } |A' - A| = 1.$$

If $A \cap B = \emptyset$, we use $A + B$ instead of $A \cup B$; if $B = \{b\}$ is a singleton, we often write $A + b$ instead of $A + \{b\}$. If $A \subset B$ we use $B - A$ instead of $B \setminus A$. We say that

$$A \text{ relates to } B \text{ (and write } A \sim B) \text{ when } A \subset B \text{ or } B \subset A;$$

if this is not the case, we write $A \not\sim B$. Note that $A \not\sim B$ if and only if two peers labelled A, B can swap chunks. The space of functions (vectors) from $\mathcal{P}(F)$ into \mathbb{Z}_+ is denoted by $\mathbb{Z}_+^{\mathcal{P}(F)}$. The stochastic model will take values in this space. The deterministic model will evolve in $\mathbb{R}_+^{\mathcal{P}(F)}$. We let $e_A \in \mathbb{Z}_+^{\mathcal{P}(F)}$ be the vector with coordinates

$$e_A^B := \mathbf{1}(A = B), \quad B \in \mathcal{P}(F).$$

For $x \in \mathbb{Z}_+^{\mathcal{P}(F)}$ or $\mathbb{R}_+^{\mathcal{P}(F)}$ we let $|x| := \sum_{A \in \mathcal{P}(F)} |x^A|$. If $\mathcal{A} \subset \mathcal{P}(F)$ then the \mathcal{A} -face $\mathbb{R}_+^{\mathcal{A}}$ of $\mathbb{R}_+^{\mathcal{P}(F)}$ is defined by $\mathbb{R}_+^{\mathcal{A}} := \{x \in \mathbb{R}_+^{\mathcal{P}(F)} : \sum_{A \in \mathcal{A}} x^A = 0, \prod_{B \notin \mathcal{A}} x^B > 0\}$.

2.3 Defining the rates of individual interactions

We follow the logic of stochastic modelling of chemical reactions or epidemics and assume that the chance of a particular interaction occurring in a short interval of time is proportional to the number of ways of selecting the peers needed for this interaction [13]. Accordingly, the interaction rates *must* be given by the formulae described below.

Consider first finding the rate of a download $A \leftarrow B$, where $A \subsetneq B$, when the state of the system is $x \in \mathbb{Z}_+^{\mathcal{P}(F)}$. There are x^A peers labelled A and x^B labelled B . We can choose them in $x^A x^B$ ways. Thus the rate of a download $A \leftarrow B$ that results into A getting *some* chunk from B should be proportional to $x^A x^B$. However, we are interested in the rate of the *specific* interaction $(A \leftarrow B) \rightsquigarrow (A', B)$, that turns A into a specific set A' differing from A by one single chunk ($A \sqsubset A'$); there are $|B - A|$ chunks that A can download from B ; the chance that picking one of them is $1/|B - A|$. Thus we have:

$$(DR) \quad \begin{cases} \text{the rate of the download } (A \leftarrow B) \rightsquigarrow (A', B) \text{ equals } \beta x^A \frac{x^B}{|B - A|}, \\ \text{as long as } A \sqsubset A' \subset B, \end{cases}$$

where $\beta > 0$.

Consider next a swap $A \rightleftharpoons B$ and assume the state is x . Picking two peers labelled A and B (provided that $A \not\sim B$) from the population is done in $x^A x^B$ ways. Thus the rate of a swap $A \rightleftharpoons B$ is proportional to $x^A x^B$. So if we *fix* two chunks $i \in A \setminus B, j \in B \setminus A$ and specify that $A' = A + j, B' = B + i$, then the chance of picking i from $A \setminus B$ and j from $B \setminus A$ is $1/|A \setminus B||B \setminus A|$. Thus,

$$(SR) \quad \begin{cases} \text{the rate of the swap } (A \rightleftharpoons B) \rightsquigarrow (A', B') \text{ equals } \gamma \frac{x^A x^B}{|A \setminus B||B \setminus A|}, \\ \text{a long as } A \sqsubset A', \quad B \sqsubset B', \quad A' - A \subset B, \quad B' - B \subset A, \end{cases}$$

where $\gamma > 0$.

2.4 Deriving the Markov chain rates

Having defined the rates of each individual interaction we can easily define rates $q(x, y)$ of a Markov chain in continuous time and state space $\mathbb{Z}_+^{\mathcal{P}(F)}$ as follows.

Define functions $\lambda_{A,A'}, \mu_{A,B} : \mathbb{R}^{\mathcal{P}(F)} \rightarrow \mathbb{R}$ by:

$$\lambda_{A,A'}(x) := \left[\beta x^A \sum_{C: C \supset A'} \frac{x^C}{|C - A|} \right] \mathbf{1}(A \sqsubset A') \quad (1a)$$

$$\mu_{A,B}(x) := \gamma \frac{x^A x^B}{|A \setminus B||B \setminus A|} \mathbf{1}(A \not\sim B). \quad (1b)$$

Consider also constants $\delta \geq 0$ and $\alpha^A \geq 0$ for $A \in \mathcal{P}(F)$, i.e., $\alpha \in \mathbb{R}_+^{\mathcal{P}(F)}$. The transition

rates of the closed conservative Markov chain are given by:

$$q(x, y) := \begin{cases} \lambda_{A,A'}(x), & \text{if } y = x - e_A + e_{A'} \\ \mu_{A,B}(x), & \text{if } \begin{cases} y = x - e_A - e_B + e_{A'} + e_{B'} \\ A \sqsubset A', B \sqsubset B', A' - A \subset B, B' - B \subset A, \end{cases} \\ \alpha^A & \text{if } y = x + e_A \\ \delta x^F & \text{if } y = x - e_F \\ 0, & \text{for any other value of } y \neq x, \end{cases} \quad (2)$$

where x ranges in $\mathbb{Z}_+^{\mathcal{P}(F)}$.

A little justification of the first two cases is needed: that $q(x, x - e_A - e_B + e_{A'} + e_{B'}) = \mu_{A,B}(x)$ is straightforward. It corresponds to a swap, which is only possible when $A \sqsubset A', B \sqsubset B', A' - A \subset B, B' - B \subset A$. The swap rate was defined by (SR). To see that $q(x, x - e_A + e_{A'}) = \lambda_{A,A'}(x)$ we observe that a peer labelled A can change its label to $A' \sqsupset A$ by downloading a chunk from some set C that contains A' , so we sum the rates (DR) over all these possible individual interactions to obtain the first line in (2). We can think of having Poisson process of arrivals of new peers at rate $|\alpha|$, and that each arriving peer is labelled A with probability $\alpha^A/|\alpha|$. Peers can depart, by definition, only when they are labelled F and it takes an exponentially distributed amount of time (with mean $1/\delta$) for a departure to occur. Thus, $q(x, x - e_F) = \delta x^F$. We shall let \mathbf{Q} denote the generator of the chain, i.e. $\mathbf{Q}f(x) = \sum_y (f(y) - f(x))q(x, y)$, when f is an appropriate functional of the state space.

Definition 1 (BITTORRENT $[x_0, n, \alpha, \beta, \gamma, \delta]$). Given $x_0 \in \mathbb{Z}_+^{\mathcal{P}(F)}$ (initial configuration), $n = |F| \in \mathbb{N}$ (number of chunks), $\alpha \in \mathbb{R}_+^{\mathcal{P}(F)}$ (arrival rates), $\beta > 0$ (download rate), $\gamma \geq 0$ (swap rate), $\delta \geq 0$ (departure rate) we let BITTORRENT $[x_0, n, \alpha, \beta, \gamma, \delta]$ be a Markov chain $(X_t, t \geq 0)$ with transition rates (2) and $X_0 = x_0$. We say that the chain (network) is open if $\alpha^A > 0$ for at least one A and $\delta > 0$; it is closed if $\alpha^A = 0$ for all A ; it is conservative if it is closed and $\delta = 0$; it is dissipative if it is closed and $\delta > 0$.

In a conservative network, we have $q(x, y) = 0$ if $|y| \neq |x|$ and so $|X_t| = |X_0|$ for all $t \geq 0$. Here, the actual state space is the simplex

$$\{x \in \mathbb{Z}_+^{\mathcal{P}(F)} : |x| = N\},$$

where $N = |X_0|$. It is easy to see that the state e_F is reachable from any other state, but all rates out of e_F are zero. Hence a conservative network has e_F as a single absorbing state.

In a dissipative network, we have $|X_t| \leq |X_0|$ for all $t \geq 0$. Here the state space is

$$\{x \in \mathbb{Z}_+^{\mathcal{P}(F)} : |x| \leq N\},$$

where $N = |X_0|$. It can be seen that a dissipative network has many absorbing points.

In an open network, there are no absorbing points. On the other hand, one may wonder if certain components can escape to infinity. This is not the case:

Lemma 1. *If $\alpha^F > 0$ then the open BITTORRENT $[x, n, \beta, \gamma, \alpha, \delta]$ is positive recurrent Markov chain.*

Proof. (sketch) If $\alpha^F > 0$, $\delta > 0$ the Markov chain is irreducible. The remainder of the proof is based on a the construction of a simple Lyapunov function:

$$V(x) := |x|,$$

for which it can be shown that there is a bounded set of states K such that

$$\sup_{x \notin K} (\mathbf{Q}V)(x) < 0.$$

Perhaps the easiest way to see this is by appealing to the stability of the corresponding ODE limit; see Theorem 1 below and [7]. \square

2.5 Example: $n = 1$

Let us take the special case where the file consists of a single chunk ($n = 1$). The state here is $x = (x^\emptyset, x^1 := x^F)$. The rates are:

$$\begin{aligned} q((x^\emptyset, x^1), (x^\emptyset + 1, x^1)) &= \alpha^\emptyset \\ q((x^\emptyset, x^1), (x^\emptyset, x^1 + 1)) &= \alpha^1 \\ q((x^\emptyset, x^1), (x^\emptyset - 1, x^1 + 1)) &= \beta x^\emptyset x^1 \\ q((x^\emptyset, x^1), (x^\emptyset, x^1 - 1)) &= \delta x^1. \end{aligned} \tag{3}$$

If $\alpha^\emptyset = \alpha^1 = 0$, this is the stochastic version of the classical (closed) Kermack-McKendrick (or susceptible-infective-removed (SIR)) model for a simple epidemic process [3]. Its absorbing points are states of the form $(x^\emptyset, 0)$. In epidemiological terminology, x^1 is the number of infected individuals, whereas x^\emptyset is the number of susceptible ones. Contrary to the epidemiological interpretation, infection *is* desirable, for infection is tantamount to downloading the file.

3 Macroscopic description: fluid limit

Analysing the Markov chain in its original form is complicated. We thus resort to a first-order approximation by an ordinary differential equation (ODE).

Let $v(x)$ be the vector field on $\mathbb{R}_+^{\mathcal{P}(F)}$ with components $v^A(x)$ defined by

$$\begin{aligned} v^A(x) &= \alpha^A - x^A (\beta \varphi_d^A(x) + \gamma \varphi_s^A(x)) \\ &\quad + \beta \sum_{B: A \subset B} \frac{\psi_d^A(x) x^B}{1 + |B \setminus A|} + \gamma \sum_{B: A \not\subset B} \frac{\psi_s^{A,B}(x) x^B}{1 + |B \setminus A|} - \delta x^F \mathbf{1}(A = F), \end{aligned} \tag{4}$$

where

$$\begin{aligned} \varphi_d^A(x) &:= \sum_{B \supset A} x^B, & \varphi_s^A(x) &:= \sum_{B \not\supset A} x^B \\ \psi_d^A(x) &:= \sum_{a \in A} x^{A-a}, & \psi_s^{A,B}(x) &:= \sum_{a \in A \cap B} x^{A-a} \end{aligned} \tag{5}$$

Consider the differential equation

$$\dot{x} = v(x) \text{ with initial condition } x_0. \quad (6)$$

Consider the sequence of stochastic models $\text{BITTORRENT}[X_{N,0}, n, N\alpha, \frac{\beta}{N}, \frac{\gamma}{N}, \delta]$ for $N \in \mathbb{N}$ and let $X_{N,t}$ be the corresponding jump Markov chain.

Theorem 1. *There is a unique smooth (analytic) solution to (6), denoted by x_t for $t \geq 0$. Also, if there is an $x_0 \in \mathbb{R}_+^{\mathcal{P}(F)}$ such that $X_{N,0}/N \rightarrow x_0$ as $N \rightarrow \infty$, then for any $T, \varepsilon > 0$,*

$$\lim_{N \rightarrow \infty} P\left(\sup_{0 \leq t \leq T} |N^{-1}X_{N,t} - x_t| > \varepsilon\right) = 0.$$

Proof. Let \mathcal{N} be the set of vectors $-e_F, e_A, -e_A + e_{A'}, -e_A - e_B + e_{A'} + e_{B'}$, where $A, B \in \mathcal{P}(F)$ and $A \sqsubset A', B \sqsubset B'$. From (2), we have that $q(x, y) = 0$ if $y - x \notin \mathcal{N}$. Introduce, for each $\zeta \in \mathcal{N}$, a unit rate Poisson process Φ_ζ on the real line, and assume that these Poisson processes are independent. Consider the Markov chain (X_t) for the $\text{BITTORRENT}[X_0, n, \alpha, \beta, \gamma, \delta]$. Its rates are of the form

$$q(x, x + \zeta) = Q_\zeta(x), \quad \zeta \in \mathcal{N}, \quad (7)$$

where $Q_\zeta(x)$ is a polynomial in 2^n variables of degree 2, and which can be read directly from (2); its coefficients depend on the parameters $\alpha, \beta, \gamma, \delta$. We can represent [13, 5] (X_t) as:

$$X_t = X_0 + \sum_{\zeta \in \mathcal{N}} \zeta \Phi_\zeta\left(\int_0^t Q_\zeta(X_s) ds\right).$$

Consider now the Markov chain $\frac{1}{N}X_{N,t}$ corresponding to $\text{BITTORRENT}[X_{N,0}, n, (N\alpha^A), \beta/N, \gamma/N, \delta]$. The transition rates for $\frac{1}{N}X_{N,t}$ are

$$q(x/N, (x + \zeta)/N) = NQ_\zeta(x/N), \quad x \in \mathbb{Z}_+^{\mathcal{P}(F)}, \quad \zeta \in \mathcal{N},$$

and 0, otherwise. Here, $Q_\zeta(x)$ is the polynomial defined through (7) and (2) and we now assume that its variables range over the reals. Therefore, $\frac{1}{N}X_{N,t}$ can be represented as

$$\frac{1}{N}X_{N,t} = \frac{1}{N}X_{N,0} + \sum_{\zeta \in \mathcal{N}} \zeta \frac{1}{N} \Phi_\zeta\left(N \int_0^t Q_\zeta\left(\frac{1}{N}X_{N,s}\right) ds\right).$$

Define x_t by the (deterministic) integral equation

$$x_t = x_0 + \sum_{\zeta \in \mathcal{N}} \zeta \int_0^t Q_\zeta(x_s) ds \quad (8)$$

and assume that it is unique for all $t \geq 0$. Fix a time horizon $T > 0$ and let

$$\begin{aligned} B &:= \max_{t \leq T} |x_t|, \\ M_\zeta &:= \max_{|x| \leq B} |Q_\zeta(x)| \\ L_\zeta &:= \sup_{\substack{|x|, |y| \leq B \\ x \neq y}} \frac{|Q_\zeta(x) - Q_\zeta(y)|}{|x - y|} \\ \tau_N &:= \inf\{t > 0 : |X_{N,t}| > NB\}. \end{aligned}$$

We then have:

$$\begin{aligned}\Delta_{N,t} := \frac{X_{N,t}}{N} - x_t &= \frac{X_{N,0}}{N} - x_0 + \sum_{\zeta \in \mathcal{N}} \zeta \left[\frac{1}{N} \Phi_{\zeta} \left(N \int_0^t Q_{\zeta}(X_{N,s}/N) ds \right) - \int_0^t Q_{\zeta}(X_{N,s}/N) ds \right] \\ &\quad + \sum_{\zeta \in \mathcal{N}} \zeta \int_0^t (Q_{\zeta}(X_{N,s}/N) - Q_{\zeta}(x_s)) ds\end{aligned}$$

Suppose that $t \leq T \wedge \tau_N$. Then, for all $s \leq t$,

$$|Q_{\zeta}(X_{N,s}/N) - Q_{\zeta}(x_s)| \leq L_{\zeta} |\Delta_{N,s}|.$$

So, if we let

$$\mathcal{E}_{N,t} := \frac{X_{N,0}}{N} - x_0 + \sum_{\zeta \in \mathcal{N}} \zeta \frac{1}{N} \left[\Phi_{\zeta} \left(N \int_0^t Q_{\zeta}(X_{N,s}/N) ds \right) - N \int_0^t Q_{\zeta}(X_{N,s}/N) ds \right],$$

we have, by the Gronwall-Bellman lemma, that

$$|\Delta_{N,t}| \leq |\mathcal{E}_{N,t}| \exp \left(t \sum_{\zeta \in \mathcal{N}} |\zeta| L_{\zeta} \right), \quad \text{if } t \leq T \wedge \tau_N.$$

Let

$$\Phi_{\zeta}^*(t) := \sup_{s \leq t} |\Phi_{\zeta}(s) - s|.$$

We recall that, as $N \rightarrow \infty$,

$$\frac{1}{N} \Phi_{\zeta}^*(Nt) \rightarrow 0 \quad \text{a.s.} \quad (9)$$

If $s \leq t \leq \tau_N$, we have $X_{N,s}/N \leq B$ (definition of τ_N) and so $Q_{\zeta}(X_{N,s}/N) \leq M_{\zeta}$, implying that

$$\sup_{t \leq T \wedge \tau_N} |\mathcal{E}_{N,t}| \leq \left| \frac{X_{N,0}}{N} - x_0 \right| + \sum_{\zeta \in \mathcal{N}} |\zeta| \frac{1}{N} \Phi_{\zeta}^*(NM_{\zeta}T)$$

which converges to zero, a.s., due to (9). Since

$$\sup_{t \leq T \wedge \tau_N} |\Delta_{N,t}| \leq \sup_{t \leq T \wedge \tau_N} |\mathcal{E}_{N,t}| \exp \left(T \sum_{\zeta \in \mathcal{N}} |\zeta| L_{\zeta} \right),$$

we have

$$\sup_{t \leq T \wedge \tau_N} |\Delta_{N,t}| \rightarrow 0, \quad \text{a.s.}$$

Now observe that

$$\begin{aligned}P(\tau_N \leq T) &\leq P\left(\sup_{t \leq T \wedge \tau_N} |X_{N,t}| > NB\right) \\ &\leq P\left(\sup_{t \leq T \wedge \tau_N} |\Delta_{N,t}| + \sup_{t \leq T \wedge \tau_N} |x_t| > B\right) \rightarrow 0.\end{aligned}$$

So we have $\sup_{t \leq T} |\Delta_{N,t}| \rightarrow 0$ a.s.

To show that x_t , defined via (8), satisfies the ODE $\dot{x} = v(x)$ with v given by (4) is a matter of straightforward (but tedious) algebra, see the Appendix.

Uniqueness and analyticity of the solution of the ODE is immediate from the form of the vector field (its components are polynomials of degree 2 and hence locally Lipschitzian).

To show that the trajectories do not explode, we consider the function

$$V(x) := \sum_A x^A.$$

It is a matter of algebra to check that

$$\langle \nabla V(x), v(x) \rangle = \sum_A v^A(x) = \sum_A \alpha^A - \delta x^F$$

which (since $\delta > 0$) is negative and bounded away from zero for x outside a bounded set of $\mathbb{R}_+^{\mathcal{P}(F)}$ containing the origin. We then apply the Lyapunov criterion for ODEs to conclude that x_t is defined for all $t \geq 0$ and this justifies the fact that we could choose an arbitrary time horizon T earlier in the proof. \square

Comment: The quantities defined in (5), have physical meanings as follows:

$$\varphi_d^A(x) := \sum_{B \supset A} x^B = \text{no. of peers from which an } A\text{-peer can download,}$$

$$\varphi_s^A(x) := \sum_{B \not\supset A} x^B = \text{no. of peers an } A\text{-peer can swap with,}$$

$$\psi_d^A(x) := \sum_{a \in A} x^{A-a} = \text{no. of peers which can assume label } A \text{ after a download,}$$

$$\psi_s^{A,B}(x) := \sum_{a \in A \cap B} x^{A-a} = \text{no. of peers which can assume label } A \text{ after a } B\text{-peer swap.}$$

It is helpful to keep these in mind because they aid in writing down the various parts of $v(x)$, again, see the Appendix.

4 Behaviour of the limiting ODE

Concerning the ODE $\dot{x} = v(x)$ we consider again three cases, just as in the stochastic model: an open system ($\alpha^A > 0$ for at least one A and $\delta > 0$), a closed dissipative system ($\alpha^A = 0$ for all A and $\delta > 0$), and a closed conservative system ($\alpha^A = 0$ for all A and $\delta = 0$). Qualitatively, the behaviour is different in each case. In this section, we will try to exemplify this behaviour by means of examples.

4.1 The ODE in the absence of BitTorrent incentives

Absence of BitTorrent incentives means that the file is not split into chunks, i.e. $n = 1$. Thus, the Markov chain is $X_t = (X_t^\emptyset, X_t^1 := X_t^F)$, i.e. two-valued $A \in \{\emptyset, 1\}$. The rates for this case were reported earlier in (3). To find the fluid limit, we use (5) and (4), keeping in mind the interpretation of each of the terms in (5).

For $A = \emptyset$, we have: the number of peers from which an \emptyset -peer can download from is $\varphi_d^\emptyset(x) = x^1$; the number of peers that can swap with an \emptyset -peer is $\varphi_s^\emptyset(x) = 0$; since no

peer can assume value \varnothing after an interaction, we have $\psi_d^\varnothing(x) = \psi_s^{\varnothing,B}(x) = 0$. Hence, in the formula for (4) for $v^\varnothing(x)$ only the first parenthesis survives:

$$v^\varnothing(x) = \alpha^\varnothing - \beta x^\varnothing x^1.$$

For $A = 1$ we have: $\varphi_d^1(x) = \varphi_s^1(x) = 0$, $\psi_d^1(x) = x^\varnothing$, $\psi_s^{1,\varnothing}(x) = 0$. Thus,

$$v^1(x) = \alpha^1 + \beta x^\varnothing x^1 - \delta x^1.$$

Here, by simply letting $x := x^\varnothing$ and $y := x^1$, the ODE is


$$\begin{aligned}\dot{x} &= \alpha^\varnothing - \beta xy \\ \dot{y} &= \alpha^1 + \beta xy - \delta y.\end{aligned}$$

If $\alpha^\varnothing = \alpha^1 = \delta = 0$ (closed conservative system), we have $x + y = \text{constant}$, say $= 1$, and

$$\dot{x} = -\beta x(1 - x),$$

the solution of which is the logistic function,

$$x_t = \frac{x_0}{x_0 + (1 - x_0)e^{\beta t}};$$



If $\alpha^\varnothing = \alpha^1 = 0$, but $\delta > 0$ (closed dissipative system) we have the classical deterministic SIR epidemic


$$\begin{aligned}\dot{x} &= -\beta xy \\ \dot{y} &= \beta xy - \delta y\end{aligned}$$

the integral curves of which can be found by solving

$$\frac{dy}{dx} = \frac{\delta}{\beta x} - 1,$$

whence it follows

$$y = (x_0 + y_0) + \frac{\delta}{\beta} \log(x/x_0) - x:$$



Assume $x_0 + y_0 \leq 1$. Notice that the integral curves are monotonic as long as they start from a point (x_0, y_0) with $x_0 > \delta/\beta$. They eventually converge to a point of the form $(x^*, 0)$ where

$$(x_0 + y_0) + \frac{\delta}{\beta} \log(x^*/x_0) - x^* = 0.$$

In other words, $(x_t, y_t) \rightarrow (x^*, 0)$ as $t \rightarrow \infty$ and, in fact, does so exponentially fast.

Notice that if the initial state is in the interior of the positive orthant then the boundary cannot be reached in finite time. This is in contrast to the corresponding stochastic model which can reach the boundary in finite time with positive probability. In fact, in the open case, it will reach the boundary in finite time with probability 1 due to (positive) recurrence. This remark is generic: it applies to any dimension.


Next, consider the open system case assuming, for simplicity, that $\alpha^\emptyset = \lambda > 0$, but $\alpha^1 = 0$. We have

$$\begin{aligned}\dot{x} &= \lambda - \beta xy \\ \dot{y} &= \beta xy - \delta y.\end{aligned}$$

Since $\delta > 0$, the Lyapunov function argument shows that the system is asymptotically stable. In fact, there is a unique asymptotic equilibrium which attracts all initial conditions (x_0, y_0) with $y_0 > 0$. This unique equilibrium is given by

$$x^* = \frac{\delta}{\beta}, \quad y^* = \frac{\lambda}{\delta},$$

as can be seen by setting the right hand side of the ODE equal to zero [3]. It should be noticed that the trajectories can be spirals around (x^*, y^*) . A typical situation is shown below.



In this vector field plot, we took $\beta = 3$, $\lambda = 5$, $\delta = 4$, so $(x^*, y^*) = (1.33, 1.25)$. If $y_0 = 0$ then the trajectory converges to infinity. Indeed, if no peers initially possess the file and no such peers ever show up, then the only thing that can happen is an accumulation, at rate λ , of ever demanding peers. The remedy is, clearly, the imposition of $\alpha^1 > 0$ no matter how small; then (x^*, y^*) is globally attracting for all initial states in the closed positive orthant.

In this last example, we find that the eigenvalues of the differential of the vector field at (x^*, y^*) are complex conjugates if and only if $\lambda\beta < 4\delta^2$, and so this is the condition for the spiralling of trajectories.

4.2 The ODE for $n = 2$ chunks

Here A can take 4 values: $\emptyset, \{1\}, \{2\}, \{1, 2\} := F$. We work out the expression for each component separately. Since an \emptyset -peer can only download from $\varphi_d^\emptyset(x) = x^1 + x^2 + x^{12}$ peers[†], but no peer can download from it or swap with it, we have $\varphi_s^\emptyset(x) = \psi_d^\emptyset(x) = \psi_s^\emptyset(x) = 0$, and so

$$v^\emptyset(x) = \alpha^\emptyset - \beta x^\emptyset(x^1 + x^2 + x^{12}).$$

A 1-peer can download from $\varphi_d^1(x) = x^{12}$ peers; it can swap with $\varphi_s^1(x) = x^2$ peers; the number of peers that take value 1 after a download is $\psi_d^1(x) + x^\emptyset$; and, since no peer can take the label 1 after a swap, we have $\psi_s^{1,B}(x) = 0 \forall B$. We thus have

$$\begin{aligned} v^1(x) &= \alpha^1 - x^1(\beta x^{12} + \gamma x^2) + \beta \psi_d^1(x) \sum_{B:1 \in B} \frac{x^B}{1 + |B \setminus \{1\}|} \\ &= \alpha^1 - x^1(\beta x^{12} + \gamma x^2) + \beta x^\emptyset(x^1 + \frac{1}{2}x^{12}). \end{aligned}$$

The expression for $v^2(x)$ is symmetric to that of $v^1(x)$. A 12-peer can neither download from or swap with anyone, so $\varphi_d^{12}(x) = \varphi_s^{12}(x) = 0$; there are $\psi_d^{12}(x) = x^1 + x^2$ peers which can take label 12 after a download; since only a 2-peer can assume label 12 after a swap, we have $\psi_s^{12,1}(x) = x^2$; and, for the same reason, $\psi_s^{12,2}(x) = x^1$. Thus,

$$\begin{aligned} v^{12}(x) &= \alpha^{12} + \beta \psi_d^{12}(x)x^{12} + \frac{\gamma \psi_s^{12,1}(x)x^1}{1 + |\{1\} \setminus \{1, 2\}|} + \frac{\gamma \psi_s^{12,2}(x)x^2}{1 + |\{2\} \setminus \{1, 2\}|} - \delta x^{12} \\ &= \alpha^{12} + \beta(x^1 + x^2)x^{12} + \gamma(x^2x^1 + x^1x^2) - \delta x^{12}. \end{aligned}$$

We thus have,

$$\begin{aligned} \dot{x}^\emptyset &= \alpha^\emptyset - \beta x^\emptyset(x^1 + x^2 + x^{12}) \\ \dot{x}^1 &= \alpha^1 - x^1(\beta x^{12} + \gamma x^2) + \beta x^\emptyset(x^1 + \frac{1}{2}x^{12}) \\ \dot{x}^2 &= \alpha^2 - x^2(\beta x^{12} + \gamma x^1) + \beta x^\emptyset(x^2 + \frac{1}{2}x^{12}) \\ \dot{x}^{12} &= \alpha^{12} + \beta(x^1 + x^2)x^{12} + 2\gamma x^1x^2 - \delta x^{12}. \end{aligned}$$

Case 1: closed conservative system. Consider the $n = 2$ case again, and assume that $\alpha^A = 0$ for all A , $\delta = 0$ (a closed conservative system). Assume further that $\gamma = 0$. Let

$$x = x^\emptyset, \quad u = x^1 + x^2, \quad w = x^{12}.$$

We see a reduction in dimension from 4 to 3 (owing to that fact that the sum of the coordinates is constant) and a further reduction from 3 to 2 (owing to the fact that, for $\gamma = 0$, the vector field depends on x^1, x^2 through their sum). We have:

$$\begin{aligned} \dot{x} &= -\beta x(u + w) \\ \dot{u} &= -\beta uw + \beta x(u + w) \\ \dot{w} &= \beta uw. \end{aligned}$$

On assuming that


$$x + u + w = 1$$

[†]Where $x^{12} = x^F$ and indexes 1 and 2 represent the chunks that form a partition of F .

and eliminating the variable w we obtain

$$\begin{aligned}\dot{x} &= -\beta x(1-x) \\ \dot{u} &= \beta u^2 - \beta u(1-x) + \beta x(1-x).\end{aligned}$$

We can qualitatively see the behaviour of this ODE by looking at the vector field in the $x - u$ plane. A typical picture is as follows:



The first is an autonomous equation, encountered earlier. Its solution is $x_t = x_0/(x_0 + (1 - x_0)^{\beta t})$. Letting

$$f := 1 - w = x + u$$

we have

$$\dot{f} = \beta f(1 - f) + \beta(1 - f)x_t$$

Assume that $w_0 > 0$, i.e. $f_0 < 1$. The solution of the last ODE is:

$$f_t = \frac{f_0(1 - f_0)^{-1} + x_0\beta t}{x_0 + (1 - x_0)e^{\beta t} + f_0(1 - f_0)^{-1} + x_0\beta t}.$$


Hence

$$w_t = \frac{x_0 + (1 - x_0)e^{\beta t}}{x_0 + (1 - x_0)e^{\beta t} + x_0\beta t + (1 - w_0)w_0^{-1}}.$$

We have $\lim_{t \rightarrow \infty} w_t = 1$, as expected. From this, we can estimate the time $\tau = \tau(x_0, w_0, \varepsilon)$ required for the system to reach the set $\{(x, u, w) : w > 1 - \varepsilon\}$, where $0 < \varepsilon < 1$ is a given (small) number. This time can be translated into an estimate for the expected time required for the original stochastic system to be absorbed by the state where all peers possess the full file. We shall not attempt the justification of this statement here. The time τ is the solution of the transcendental equation

$$\frac{1 - w_0}{w_0} + x_0\beta\tau = \frac{\varepsilon}{1 - \varepsilon}(x_0 + (1 - x_0)e^{\beta\tau}).$$

Typical behaviour of this time as a function of x_0 is as follows:



The three graphs correspond to varying values of β ranging from 1 (top curve) to 5 (bottom curve) and to $\varepsilon = 0.001$, $w_0 = 0.1$.

Case 2: closed dissipative system. Consider again the $n = 2$ case with no arrivals, with swap rate γ equal to zero, but with departure rate $\delta > 0$. Using the same notation as above, we have

$$\begin{aligned}\dot{x} &= -\beta x(u + w) \\ \dot{u} &= -\beta uw + \beta x(u + w) \\ \dot{w} &= \beta uw - \delta w.\end{aligned}$$

To substitute out one parameter, change the time variable to $s = \beta t$, let $\rho = \delta/\beta$, and write x' for dx/ds , so that

$$\begin{aligned}x' &= -x(u + w) \\ u' &= -uw + x(u + w) \\ w' &= uw - \rho w.\end{aligned}$$

Assume

$$x_0 + u_0 + w_0 = 1,$$

so $x_t + u_t + w_t < 1$ for all $t > 0$. We cannot eliminate the variable w now since there is no obvious conserved quantity, but we can study the equilibria of the system. It seems that the ODE has no obvious analytical solution. However it can either be integrated numerically or its solution in terms of power series can be easily found. Setting the right hand side equal to zero we see that, necessarily, $w = 0$, which leaves us with the possibility $xu = 0$. So points of the form

$$(x, u, w) = (x, 0, 0), (0, u, 0)$$

are equilibria. But not all of them are stable. For example, any point of the form $(0, u, 0)$ with $u > \rho$ is an unstable equilibrium. Specifically, the differential of the vector field is given by

$$Dv = \begin{pmatrix} -(u + w) & -x & -x \\ u + w & x - w & x - u \\ 0 & w & u - \rho \end{pmatrix}$$

Evaluated at $(0, u, 0)$, it gives $Dv = \begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & -u \\ 0 & 0 & u - \rho \end{pmatrix}$, and this has eigenvalues $0, -u, u - \rho$.

Evaluated at $(x, 0, 0)$, it gives $Dv = \begin{pmatrix} 0 & -x & -x \\ 0 & x & x \\ 0 & 0 & \rho \end{pmatrix}$, and this has eigenvalues $0, x, -\rho$.

Therefore, the only stable equilibria are points of the form

$$(x, u, w) = (0, u, 0), \quad 0 \leq u < \rho.$$

In terms of the original variables, the stable equilibria are

$$(x^\emptyset, x^1, x^2, x^{12}) = (0, x^1, x^2, 0), \quad 0 \leq x^1 + x^2 < \rho.$$

This is as expected: since there is no swapping ($\gamma = 0$), the system eventually settles to a situation where there are peers with label 1 and peers with label 2. Had γ been positive, x^1, x^2 could not have simultaneously been positive in equilibrium.


Case 3: Open system. Consider the situation as in Case 2, but add arrivals of peers (known as seeds) possessing the full file. Choosing variables appropriately, we have

$$\begin{aligned} x' &= -x(u + w) \\ u' &= -uw + x(u + w) \\ w' &= \lambda + uw - \rho w. \end{aligned}$$

In terms of the original variables, we here have $\alpha^\emptyset = \alpha^1 = \alpha^2 = 0$, $\alpha^{12} = \lambda\beta > 0$. Here we see that the system eventually settles to the stable equilibrium

$$(x, u, w) = (0, 0, \lambda/\rho).$$

We can easily see that the eigenvalues of the differential of the vector field at the stable equilibrium are both real and negative: $-\lambda/\rho$ and $-\rho$ (the first one has algebraic multiplicity 1 but geometric multiplicity 2). Hence there is no possibility of spiralling.



4.3 Time to settle

Estimating the time for the system to reach an equilibrium requires cooking up an appropriate Lyapunov function. The obvious Lyapunov function used earlier gives a crude lower bound.

Consider a general deterministic system (open or closed), i.e. the differential equation $\dot{x} = v(x)$ with $v(x)$ given by (4). Let $V(x) = |x| = \sum_{A \subset F} x^A$. Then

$$\dot{V} = |\alpha| - \delta x^F,$$

where $|\alpha| = \sum_{A \subset F} \alpha^A$ is the total arrival rate. Using the naive inequality $x^F \leq V$ we obtain the following:

Corollary 1. *If the system starts from x_0 and if $x_t \rightarrow x^*$ then the time τ_r required for the trajectory to reach an r -ball centred at x^* satisfies*

$$\tau_r \geq \frac{1}{\delta} \log \frac{|x_0| + |\alpha|}{|x^*| + r}.$$

In the other direction, consider the last component $v^F(x)$ of the vector field. Since no seed (F -peer) can download or swap, we have $\varphi_d^F(x) = \varphi_s^F(x) = 0$. So

$$\begin{aligned} v^F(x) &= \alpha^F + \beta \psi_d^F(x) x^F + \gamma \sum_{B \subset F} \psi_s^{F,B}(x) x^B - \delta x^F \\ &= \alpha^F + \beta \sum_{j \in F} x^F x^{F-j} + \gamma \sum_{B \subset F} \sum_{i \in B} x^B x^{F-i} - \delta x^F \\ &=: \alpha^F + v_+^F(x) - \delta x^F, \end{aligned}$$

where $v_+^F(x)$ is a quadratic form with positive coefficients.

Assume that the system is closed, so that $\alpha^F = 0$ in particular. Let $\overline{v_+^F}$ be an upper bound on $v_+^F(x)$, i.e., a (good) bound that depends on the initial state x_0 and the parameters β, γ .

Corollary 2. *If the system is closed and $\overline{v_+^F} < \delta$, then*

$$\tau_r \leq \frac{1}{\overline{v_+^F} - \delta} \log \frac{x_0^F}{r}.$$

We conjecture that $v_+^F(x_t)$ decreases along the trajectory x_t as long as $v_+^F(x_0) < \delta$. If so, the bound above is $\frac{1}{v_+^F(x_0) - \delta} \log \frac{x_0^F}{r}$.

5 An example of the evaluation of performance improvement in presence of BitTorrent incentives

We address the following question: When is it advantageous to split a file into chunks? In other words, assuming we fix certain system parameters (e.g., arrival rates), will peers acquire the file faster if the file is split into chunks? We attempt here to answer the question in a simple case only by using the deterministic approximation. Let λ be the total peer

arrival rate. Let β be the download rate. Assume that only \emptyset peers arrive exogenously. In the absence of BitTorrent incentives, we have the single-chunk case

$$\begin{aligned}\dot{x}^\emptyset &= \lambda - \beta x^\emptyset x^1 \\ \dot{x}^1 &= \beta x^\emptyset x^1 - \delta x^1.\end{aligned}$$

The globally attracting stable equilibrium is given by

$$x^* = (\delta/\beta, \lambda/\delta).$$

Consider splitting into $n = 2$ chunks. Let \tilde{x} be the state of the system. Suppose that the new parameters are $\tilde{\lambda} = \lambda$, $\tilde{\delta} = \delta$, $\tilde{\beta}$, $\tilde{\gamma}$. Then

$$\begin{aligned}\dot{\tilde{x}}^\emptyset &= \lambda - \tilde{\beta}\tilde{x}^\emptyset(\tilde{x}^1 + \tilde{x}^2 + \tilde{x}^{12}) \\ \dot{\tilde{x}}^1 &= -\tilde{x}^1(\tilde{\beta}\tilde{x}^{12} + \tilde{\gamma}\tilde{x}^2) + \tilde{\beta}\tilde{x}^\emptyset(\tilde{x}^1 + \frac{1}{2}\tilde{x}^{12}) \\ \dot{\tilde{x}}^2 &= -\tilde{x}^2(\tilde{\beta}\tilde{x}^{12} + \tilde{\gamma}\tilde{x}^1) + \tilde{\beta}\tilde{x}^\emptyset(\tilde{x}^2 + \frac{1}{2}\tilde{x}^{12}) \\ \dot{\tilde{x}}^{12} &= \tilde{\beta}(\tilde{x}^1 + \tilde{x}^2)\tilde{x}^{12} + 2\tilde{\gamma}\tilde{x}^1\tilde{x}^2 - \delta\tilde{x}^{12}.\end{aligned}$$

The new equilibrium is easily found to be

$$\tilde{x}^* = \left(\frac{\delta}{\tilde{\beta}} \left(\frac{\delta}{\tilde{\lambda}} u + 1 \right)^{-1}, \frac{u}{2}, \frac{u}{2}, \frac{\lambda}{\delta} \right),$$

where u is the positive number which solves

$$q(u) = 0,$$

and

$$q(u) := u^2 + \frac{2\tilde{\beta}\lambda}{\tilde{\gamma}\delta}u - \frac{2\lambda}{\tilde{\gamma}}. \quad (10)$$

To see this, set the vector field equal to zero and solve for \tilde{x}^A , $A \in \{\emptyset, 1, 2, 12 := F\}$. The simplest way to obtain the solution is by first adding the equations; this gives

$$\lambda - \delta\tilde{x}^{12} = 0,$$

whence $\tilde{x}^{12} = \lambda/\delta$. Then add the middle two equations after setting $u = \tilde{x}^1 + \tilde{x}^2$:

$$-\tilde{\beta}u\tilde{x}^{12} - 2\tilde{\gamma}\tilde{x}^1\tilde{x}^2 + \tilde{\beta}\tilde{x}^\emptyset(u + \tilde{x}^{12}) = 0.$$

Replace \tilde{x}^{12} by λ/δ and observe that, due to symmetry, $\tilde{x}^1 = \tilde{x}^2 = u/2$. This gives the quadratic equation $q(u) = 0$ with q defined by (10). Finally, the first equation becomes

$$\lambda - \tilde{\beta}\tilde{x}^\emptyset(u + \lambda/\delta) = 0,$$

which is solved for \tilde{x}^\emptyset giving:

$$\tilde{x}^{*\emptyset} = \frac{\lambda}{\tilde{\beta}} \left(u + \frac{\lambda}{\delta} \right)^{-1} = \frac{\delta}{\tilde{\beta}} \left(\frac{\delta}{\lambda} u + 1 \right)^{-1} < \frac{\delta}{\tilde{\beta}}.$$

Thus:

Corollary 3. *If $\tilde{\beta} \geq \beta$ then*

$$\tilde{x}^{*\emptyset} < x^{*\emptyset}.$$

So, by introducing splitting into chunks, we have fewer peers who have no parts of the file at all. Using Little's theorem (see below), this can be translated into smaller waiting time from the time a peer arrives until he gets his first chunk.

Suppose now we are interested in determining how long it will take for a newly arrived peer to acquire the full file. On the average, a peer spends time equal to $\lambda^{-1}|x^*|$ before it exits the system. During last part of his sojourn interval (which is a random variable with mean $1/\delta$), the peer possess the full file. It thus takes on the average $\lambda^{-1}|x^*| - \delta^{-1}$ for a peer to acquire the full file. Since we assume that $\tilde{\lambda} = \lambda$, $\tilde{\delta} = \delta$, it suffices to show that

$$|x^*| > |\tilde{x}^*|.$$

But

$$\begin{aligned} |x^*| - |\tilde{x}^*| &= \left[\frac{\delta}{\beta} + \frac{\lambda}{\delta} \right] - \left[\frac{\delta}{\tilde{\beta}} \left(\frac{\delta}{\lambda} u + 1 \right)^{-1} + u + \frac{\lambda}{\delta} \right] \\ &= \frac{\delta}{\beta} - \frac{\delta}{\tilde{\beta}} \left(\frac{\delta}{\lambda} u + 1 \right)^{-1} - u - \frac{\lambda}{\delta} \\ &= \left(\frac{\delta}{\lambda} u + 1 \right) \left[\left(\frac{\delta}{\beta} - \frac{\delta}{\tilde{\beta}} \right) + \left(\frac{\delta^2}{\beta\lambda} - 1 \right) u - \frac{\delta}{\lambda} u^2 \right] \end{aligned}$$

recalling $u > 0$ solves $q(u) = 0$. So, $|x^*| - |\tilde{x}^*| > 0$ if and only if

$$0 > \tilde{q}(u) := u^2 - \left(\frac{\delta}{\beta} - \frac{\lambda}{\delta} \right) u - \left(\frac{\lambda}{\beta} - \frac{\lambda}{\tilde{\beta}} \right). \quad (11)$$

Define \tilde{u} as the unique positive number which satisfies

$$\tilde{q}(\tilde{u}) = 0.$$

Corollary 4. *If $\tilde{\beta} \geq \beta$, a necessary and sufficient condition for $|x^*| > |\tilde{x}^*|$ is $u < \tilde{u}$.*

This gives a set of non-vacuous conditions for achieving improvement of performance by the introduction of BitTorrent incentives.[‡] It can be proved that if the parameters β , δ , $\tilde{\beta}$, $\tilde{\gamma}$ are fixed and if $\tilde{\beta} \geq \beta$ then there exists a value λ_0 such that for all $\lambda < \lambda_0$ the inequality $u < \tilde{u}$ holds. To prove this, we observe that

$$u < \tilde{u} \iff q(\tilde{u}) > 0$$

and study the behaviour of $q(\tilde{u})$ as a function of λ in a neighbourhood of zero.

We conjecture that an algebraic condition involving quadratics like q and \tilde{q} is valid for larger values of n also.

To justify the use of deterministic approximation for estimating performance measures, and, specifically, the use of mean values, we need to show that as $N \rightarrow \infty$, we can approximate stationary averages in the original stochastic network by equilibria of the resulting

[‡]The inequality conditions were mistakenly reversed in the corresponding results of [11].

ODE. It is easy to show that the a.s. convergence to the ODE limit can be translated into convergence of the means, using a uniform integrability argument. Namely,

$$\frac{1}{N}EX_{N,t} \xrightarrow[N \rightarrow \infty]{} x_t \xrightarrow[t \rightarrow \infty]{} x^*,$$

where the second limit concerns the behaviour of the ODE alone. On the other hand, if we fix N and look at the asymptotic behaviour of the process $\frac{1}{N}X_{N,t}$ as $t \rightarrow \infty$, we have

$$\frac{1}{N}EX_{N,t} \xrightarrow[t \rightarrow \infty]{} \frac{1}{N}E\tilde{X}_N,$$

where the law of \tilde{X}_N is the stationary distribution of the chain $(\frac{1}{N}X_{N,t})_{t \geq 0}$. It can be proved that $\frac{1}{N}E\tilde{X}_N \rightarrow x^*$, as $N \rightarrow \infty$. Arguments for this will be considered in future work. More detailed estimates on the discrepancy between the stochastic and deterministic systems can be found in the recent survey paper [6].

We can also explain the use of $|x^*|$ as a measure of the sojourn time in the system of a peer, by first using the approximation outlined above and then appealing to Little's law. This is as follows.

Consider an open BITTORRENT $[X_0, n, \alpha, \beta, \gamma, \delta]$, i.e. $|\alpha| > 0$, $\delta > 0$. We know that the Markov chain (X_t) is positive recurrent and has thus a unique stationary distribution. It makes sense to assess the performance of the network by looking at steady-state performance measures, such as the mean time it takes for an \emptyset -peer to become an F -peer (a seed). Consider then the process $(\tilde{X}_t, t \in \mathbb{R})$ defined to be a stationary Markov process with time index \mathbb{R} and transition rates as those of (X_t) . The law of the process $(\tilde{X}_t, t \in \mathbb{R})$ is unique. Let $T_k^A, k \in \mathbb{Z}$ be the times at which A -peers arrive (and, say, $T_0^A \leq 0 < T_1^A$, by convention). These are the points of a stationary Poisson process in \mathbb{R} with rate α^A . Let W_k^A be the sojourn time in the system of a peer arriving at time T_k^A . Since, by assumption, a peer departs only after it has acquired the full set, the time W_k^A is the sum of the times it takes for the peer to become a seed plus the time that the peer hangs out in the system after becoming a seed (the latter is an exponential time with mean $1/\delta$). Clearly then, for all $t \in \mathbb{R}$,

$$\sum_{B \supset A} \tilde{X}_t^B = \sum_{k \in \mathbb{Z}} \mathbf{1}(T_k^A \leq t < T_k^A + W_k^A).$$

Using Campbell's formula, we obtain

$$\sum_{B \supset A} E\tilde{X}_0^B = \alpha^A E^A W_0^A, \quad (12)$$

where E^A is expectation with respect to P^A —the Palm probability of P with respect to the point process $(T_k^A, k \in \mathbb{Z})$.

In particular, with $A = \emptyset$, and $\lambda = \alpha^\emptyset$, we have that

$$E^\emptyset W_0^\emptyset = \frac{1}{\lambda} E|\tilde{X}_0|,$$

which can be read as: the mean sojourn time of a \emptyset -peer is, in steady state, equal to the mean number of peers in the system divided by the rate of arrivals of \emptyset -peers. If N is a parameter of the process as in Theorem 1 then, λ being proportional to N , we have that the right hand side converges to something that is proportional to $|x^*|$, as required.

6 Diffusion approximation

Using the functional central limit theorem for Poisson processes, we can prove, by standard methods, the following: Again consider the sequence $\text{BITTORRENT}[X_{N,0}, n, N\alpha, \frac{\beta}{N}, \frac{\gamma}{N}, \delta]$ for $N \in \mathbb{N}$, and let $X_{N,t}$ be the corresponding jump Markov chain. Let $(x_t, t \geq 0)$ be the solution to the ODE $\dot{x} = v(x)$ with initial condition x_0 . Let

$$Y_{N,t} := \sqrt{N}(X_{N,t}/N - x_t).$$

Let W_ζ , $\zeta \in \mathcal{N}$, be i.i.d. standard Brownian motions in \mathbb{R} . Finally, define the (time-inhomogeneous) Gaussian diffusion process Y by

$$dY_t = \sum_{\zeta \in \mathcal{N}} \zeta \sqrt{Q_\zeta(x_t)} dW_{\zeta,t} + Dv(x_t) Y_t dt,$$

where $Dv(x)$ is the matrix of partial derivatives of $v(x)$.

Theorem 2. *If $\sqrt{N}(X_{N,t}/N - x_0) \rightarrow 0$ as $N \rightarrow \infty$, where $x_0 \in \mathbb{R}_+^{\mathcal{P}(F)}$, then the law of Y_N (as a sequence of probability measures in $D[0, \infty)$ with the topology of uniform convergence on compacta) converges weakly to the law of Y .*

The proof of this theorem is omitted but the reader is referred to [13] for the relevant arguments.

7 Final remarks, open problems and future work

7.1 Rates of convergence

We can obtain a computable rate of convergence of the stochastic model to the ODE by using a combination of large deviations techniques with the solution of two optimisation problems. The idea is basically implicit in the proof of Theorem 1 and this is the reason we wrote the proof explicitly in terms of the driving Poisson processes Φ_ζ .

The first problem is so that we obtain an estimate of the maximum value of M_ζ of $Q_\zeta(x)$. In the case of a closed network, this is a quadratic optimisation problem over the polyhedron $\{|x| \leq 1\}$.

The second optimisation problem is for an estimate for L_ζ , which can be translated to an estimate for the norm of the gradient $\nabla Q_\zeta(x)$. In the case of a closed network, we can estimate this by solving a (large) number of linear programming problems. Of course, only estimates are needed.

7.2 Conjectures

The first one concerns the behaviour of $v_+^F(x_t)$, and was stated at the end of Section 4. The second concerning generalisation of Corollary 4, was stated in Section 5.

The third conjecture is more vague: it basically says that we can evaluate the performance improvement by using a large number of chunks (say 100), by solving a number of

quadratic inequalities. To this end, it should be remarked that in a deterministic open network, the unique equilibrium x^* can be found by solving $n + 1$ equations in $n + 1$ unknowns, provided that the rate of arrivals of A -peers depends on A through its cardinality alone:

$$\alpha^A = \lambda_k \text{ if } |A| = k \text{ chunks..}$$

Indeed, by symmetry of the vector field, we see that

$$x^{*A} = x^{*B} \text{ if } |A| = |B|.$$

So if we define

$$z^k := \sum_{|A|=k} x^{*A}$$

we will have

$$x^{*A} = \binom{n}{k}^{-1} z^k, \text{ if } |A| = k.$$

Hence if we let

$$V^k(z^0, \dots, z^n) := \sum_{|A|=k} v^A(x^*),$$

where z and x^* are related as above, then the equations we need to solve are

$$V^k(z^0, \dots, z^n) = 0, \quad k = 0, \dots, n.$$

7.3 A reduction of dimension for a balanced ODE

If we are interested not only in the equilibria but also in a more detailed study of the transient behaviour of the ODE, then we can obtain a rough idea (and bounds) by making the assumption of full symmetry, i.e., we assume that the arrival rates α^A and the initial states x_0^A depend on A only through $|A|$. Then the trajectory itself x_t^A depends only on the cardinality of A and so we can reduce the ODE to an $(n + 1)$ -dimensional one. Such a symmetrised ODE can yield more detailed information on the time to reach a small neighbourhood of the equilibrium point. (Note that our bounds in §4.3 are very crude.)

7.4 Non-Poissonian assumptions

It may be more reasonable in practise to assume that the time it takes for a chunk to be downloaded or swapped is a random variable with a heavy-tailed distribution. This is not captured by our model. Indeed, the interaction times are not part of our model at all. A new, more detailed, model should be worked out.

However, a crude capture of this phenomenon is the replacement of the Poisson processes Φ_ζ by more general point processes, perhaps with heavy-tailed inter-event times. As long as these processes obey a functional law of large numbers, we can (by possibly modifying the scaling parameters) rephrase Theorem 1 and repeat the proof in this more general case.

References

- [1] BitTorrent: www.bittorrent.com
- [2] B. Cohen. Incentives Build Robustness in BitTorrent. *Workshop on Economics of Peer-to-Peer Systems*, Berkeley, CA, USA, May 2003.
- [3] D.J. Daley and J. Gani. *Epidemic Modeling, an Introduction*. Cambridge University Press, 1999.
- [4] DC Info (Distributed Systems Newsletter). Vol. XXII, Issue 8, June 30, 2008; Available at: <http://www.dcia.info/news/#Newsletters>
- [5] S.N. Ethier and T.G. Kurtz. *Markov Processes: Characterization and Convergence*. Wiley, New, York, 1986.
- [6] R.W.R. Darling and J.R. Norris. Differential equation approximations for Markov chains. *Probability Surveys* **5**, 37-79, 2008.
- [7] S. Foss and T. Konstantopoulos. An overview of some stochastic stability methods. *Journal of the Operations Research Society of Japan* **47** 275-303, 2004.
- [8] Z. Ge, D.R. Figueiredo, S. Jaiswal, J. Kurose, and D. Towsley. Modeling peer-to-peer file sharing systems. In *Proc. IEEE INFOCOM*, San Francisco, April 2004.
- [9] L. Guo, S. Chen, Z. Xiao, E. Tan, X. Ding and X. Zhang, Measurements, modeling and analysis of BitTorrent-like systems. In *Proc. Internet Measurement Conference (IMC)*, Oct. 2005.
- [10] G. Kesidis, Y. Jin, B. Mortazavi, and T. Konstantopoulos. An epidemiological model for file-sharing with BitTorrent-like incentives. In *Proc. IEEE GLOBECOM*, Nov. 2006.
- [11] G. Kesidis, T. Konstantopoulos and P. Sousi. Modeling file-sharing with BitTorrent-like incentives. In *Proc. IEEE ICASSP*, Honolulu, April 2007.
- [12] G. Kesidis, M. Vojnovic, I. Hamadeh, Y. Jin, S. Jiwasurat. A Model of the Spread of Randomly Scanning Internet Worms that Saturate Access Links. *ACM TOMACS*, May 2008.
- [13] T. Kurtz. *Approximation of Population Processes*. SIAM, Philadelphia, 1981.
- [14] A. Legout, N. Liogkas, E. Kohler and L. Zhang. Clustering and sharing incentives in BitTorrent systems. *ACM SIGMETRICS Performance Evaluation Review*, **35**(1), June 2007.
- [15] L. Massoulié and M. Vojnovic. Coupon replication systems. In *Proc. ACM SIGMETRICS*, Banff, Alberta, Canada, 2005.
- [16] P. Maymounkov and D. Mazieres. Kademlia: a peer-to-peer information system based on the XOR metric. *Proc. IPTPS*, Cambridge, MA, USA, March 2002.
- [17] B. Mortazavi and G. Kesidis. A peer-to-peer content-distribution game with a reputation-based incentive mechanism. In *Proc. IEEE Workshop on Information Theory and its Applications (ITA)*, UC San Diego, Feb. 2006.

- [18] D. Qiu and R. Srikant. Modeling and performance analysis of BitTorrent-like peer-to-peer networks. In *Proc. ACM SIGCOMM*, Portland, Oregon, 2004.
- [19] B. Turner. Generalizing BitTorrent: how to build data exchange markets (and profit from them!), Jan. 2005;
available at: <http://www.fractalscape.org/GeneralizingBitTorrent.htm>
- [20] K. Walsh and E.G. Sirer. Fighting peer-to-peer SPAM and decoys with object reputation. In *Proc. Workshop on Economics of Peer-to-Peer Systems (p2pecon)*, Philadelphia, PA, August 2005.
- [21] X. Yang and G. de Veciana. Service capacity of peer to peer networks. In *Proc. IEEE INFOCOM*, San Francisco, 2004.

A Drift calculation

We consider the set of vectors

$$\begin{aligned}\mathcal{N} = \{-e_F\} \cup \{e_A : A \subset F\} \cup \{-e_A + e_{A'} : A \subset A' \subset F\} \\ \cup \{-e_A - e_B + e_{A'} + e_{B'} : A \subset A' \subset F, B \subset B' \subset F, A' - A \subset B, B' - B \subset A\}.\end{aligned}$$

For each $\zeta \in \mathcal{N}$ we define a polynomial $Q_\zeta(x)$, by comparing (7) and (2):

$$\begin{aligned}Q_{e_A}(x) &:= \alpha^A \\ Q_{-e_F}(x) &:= \delta x^F \\ Q_{-e_A+e_{A'}}(x) &:= \lambda_{A,A'}(x) \\ Q_{-e_A-e_B+e_{A'}+e_{B'}}(x) &:= \mu_{A,B}(x)\delta_{A,A',B,B'},\end{aligned}\tag{13}$$

where $\lambda_{A,A'}(x)$, $\mu_{A,B}(x)$ are given by (1a), (1b), respectively, and

$$\delta_{A,A',B,B'} := \mathbf{1}(A \sqsubset A', A' - A \subset B, B \sqsubset B', B' - B \subset A).$$

The variable x ranges in $\mathbb{Z}_+^{\mathcal{P}(F)}$ or in $\mathbb{R}_+^{\mathcal{P}(F)}$. The algebra is the same in both cases. Define the drift vector field by $\sum_y (y - x)q(x, y)$. Comparing (13) and (2) we have

$$\sum_y (y - x)q(x, y) = \sum_{\zeta \in \mathcal{N}} \zeta Q_\zeta(x).$$

The latter sum appears in (8), in the course of the proof of Theorem 1. We shall verify that

$$u(x) := \sum_{\zeta \in \mathcal{N}} \zeta Q_\zeta(x) = v(x),$$

where $v(x)$ is defined by (4).

Consider the terms in the summation $u(x) = \sum_{\zeta \in \mathcal{N}} \zeta Q_\zeta(x)$ involving $\zeta = -e_A - e_B + e_{A'} + e_{B'}$. Notice that swapping A with B or A' with B' will not change the value of $x - e_A, x - e_B + e_{A'} + e_{B'}$, so we need to make sure to take into account this change only once in the summation. If we *simultaneously* swap A with B and A' with B' then neither

$x - e_A, x - e_B + e_{A'} + e_B$ nor the value of $Q_{-e_A-e_B+e_{A'}+e_{B'}}(x) = \mu_{A,B}(x)\delta_{A,A',B,B'}$ will change because, clearly,

$$\mu_{A,B}(x)\delta_{A,A',B,B'} = \mu_{B,A}(x)\delta_{B,B',A,A'},$$

as readily follows from (1a) and (1b). We now see that to swap A with B *without* swapping A' with B' is impossible (unless $A' = B'$). Indeed, it is an easy exercise that

$$\delta_{A,B,A',B'} = \delta_{B,A,A',B'} \Rightarrow A' = B'.$$

Taking into account this, we write

$$\begin{aligned} u(x) = \alpha^A e_A - \delta x^F e_F + \sum_{A,A'} (-e_A + e_{A'}) \lambda_{A,A'}(x) \\ + \frac{1}{2} \sum_{A,B,A',B'} (-e_A - e_B + e_{A'} + e_{B'}) \mu_{A,B}(x) \delta_{A,B,A',B'}, \end{aligned} \quad (14)$$

where the $1/2$ appears because each term must be counted exactly once. The variables A, A', B, B' in both summations are free to move over $\mathcal{P}_n(F)$ (but notice that restrictions have effectively been pushed in the definitions of $\lambda_{A,A'}$, $\mu_{A,B}$, and $\delta_{A,B,A',B'}$).

Since

$$\begin{aligned} \sum_{A,B,A',B'} e_A \mu_{A,B}(x) \delta_{A,B,A',B'} &= \sum_{A,B,A',B'} e_B \mu_{A,B}(x) \delta_{A,B,A',B'}, \\ \sum_{A,B,A',B'} e_{A'} \mu_{A,B}(x) \delta_{A,B,A',B'} &= \sum_{A,B,A',B'} e_{B'} \mu_{A,B}(x) \delta_{A,B,A',B'}, \end{aligned}$$

we have

$$\begin{aligned} v(x) = \alpha^A e_A - \delta x^F e_F - \sum_{A,A'} e_A \lambda_{A,A'}(x) + \sum_{A,A'} e_{A'} \lambda_{A,A'}(x) \\ - \sum_{A,B,A',B'} e_A \mu_{A,B}(x) \delta_{A,B,A',B'} + \sum_{A,B,A',B'} e_{A'} \mu_{A,B}(x) \delta_{A,B,A',B'}. \end{aligned}$$

Call the four sums appearing in this display as $u_I(x), u_{II}(x), u_{III}(x), u_{IV}(x)$, in this order. We use the definitions (1a), (1b) of $\lambda_{A,A'}$, $\mu_{A,B}$ and find the components of the vectors u_I, \dots, u_{IV} by hitting each one with a unit vector e_G , i.e. by taking the inner products $v_I^G = \langle e_G, u_I \rangle, \dots, v_{IV}^G = \langle e_G, u_{IV} \rangle$. We have:

$$\begin{aligned} u_I^G(x) &= - \sum_{A'} \lambda_{G,A'}(x) = - \sum_{A'} \beta x^G \sum_{B:B \supset A'} \frac{x^B}{|B-G|} \mathbf{1}(G \sqsubset A') \\ &= - \beta x^G \sum_B \frac{x^B}{|B-G|} \sum_{A'} \mathbf{1}(G \sqsubset A' \subset B) = - \beta x^G \sum_{B \supset G} x^B, \end{aligned} \quad (15)$$

where, in deriving the last equality we just observed that the number of sets A' that contain one more element than G and are contained in B is equal to $|B-G|$, as long as $G \subset B$:

$$\sum_{A'} \mathbf{1}(G \sqsubset A' \subset B) = |B-G| \mathbf{1}(G \subset B).$$

Next,

$$\begin{aligned} u_{\text{II}}^G(x) &= \sum_A \lambda_{A,G}(x) = \sum_A \beta x^A \sum_{B \supset G} \frac{x^B}{|B-A|} \mathbf{1}(A \sqsubset G) \\ &= \beta \sum_{B \supset G} x^B \sum_A \frac{x^A}{|B-A|} \mathbf{1}(A \sqsubset G) \end{aligned}$$

Notice that, in the last summation, G contains exactly one more element than A and is strictly contained in B , so $|B-A| = |B-G| + 1$. Hence

$$u_{\text{II}}^G(x) = \beta \sum_{B: B \supset G} \frac{x^B}{|B-G|+1} \sum_A x^A \mathbf{1}(A \sqsubset G) = \beta \sum_{B: B \supset G} \frac{x^B}{|B-G|+1} \sum_{g \in G} x^{G-g}. \quad (16)$$

For $u_{\text{III}}(x)$, we have:

$$\begin{aligned} u_{\text{III}}^G(x) &= - \sum_{B, A', B'} \mu_{G,B}(x) \delta_{G, A', B, B'} \\ &= -\gamma \sum_B \frac{x^G x^B}{|G \setminus B| |B \setminus G|} \cdot \sum_{A'} \mathbf{1}(G \sqsubset A', A' - G \subset B) \cdot \sum_{B'} \mathbf{1}(B \sqsubset B', B' - B \subset G) \\ &= -\gamma \sum_B \frac{x^G x^B}{|G \setminus B| |B \setminus G|} \cdot |B \setminus G| \mathbf{1}(B \setminus G \neq \emptyset) \cdot |G \setminus B| \mathbf{1}(G \setminus B \neq \emptyset) \\ &= -\gamma x^G \sum_{B \not\supset G} x^B. \quad (17) \end{aligned}$$

As for the last term, we have:

$$\begin{aligned} u_{\text{IV}}^G(x) &= \sum_{A, B, B'} \mu_{A,B}(x) \delta_{A, G, B, B'} \\ &= \gamma \sum_B \sum_A \frac{x^A x^B}{|A \setminus B| |B \setminus A|} \mathbf{1}(A \sqsubset G, G - A \subset B) \sum_{B'} \mathbf{1}(B \sqsubset B', B' - B \subset A) \\ &= \gamma \sum_B \sum_A \frac{x^A x^B}{|B \setminus A|} \mathbf{1}(A \sqsubset G, G - A \subset B) \mathbf{1}(A \setminus B \neq \emptyset) \\ &= \gamma \sum_B \frac{x^B}{|B \setminus G| + 1} \sum_A x^A \mathbf{1}(A \sqsubset G, G - A \subset B) \mathbf{1}(G \not\subset B) \\ &= \gamma \sum_B \frac{x^B}{|B \setminus G| + 1} \sum_{g \in G \cap B} x^{G-g} \mathbf{1}(G \not\subset B) \quad (18) \end{aligned}$$

Adding (15) and (17) we obtain the first part of (4), while (17) and (18) give the second part.